Electron. Commun. Probab. **20** (2015), no. 29, 1–8.
DOI: 10.1214/ECP.v20-3521
ISSN: 1083-589X

ELECTRONIC COMMUNICATIONS in PROBABILITY

arXiv:1405.2552v3 [math.PR] 2 Apr 2015

Characterisation of gradient flows on finite state Markov chains*

Helge Dietert†

Abstract

In his 2011 work, Maas has shown that the law of any time-reversible continuous-time Markov chain with finite state space evolves like a gradient flow of the relative entropy with respect to its stationary distribution. In this work we show the converse to the above by showing that if the relative law of a Markov chain with finite state space evolves like a gradient flow of the relative entropy functional, it must be time-reversible. When we allow general functionals in place of the relative entropy, we show that the law of a Markov chain evolves as gradient flow if and only if the generator of the Markov chain is real diagonalisable. Finally, we discuss what aspects of the functional are uniquely determined by the Markov chain.

Keywords: Gradient flows; Finite state Markov chains; Time-reversibility.
AMS MSC 2010: 60J27.
Submitted to ECP on May 12, 2014, final version accepted on March 1, 2015.
Supersedes arXiv:1405.2552v2.

1 Introduction

The seminal paper of Jordan, Kinderlehrer, and Otto [3] identified Markov processes in the continuous setting as gradient flows of the entropy using the Wasserstein distance. This understanding lead to many new results (see Villani [6] for an overview). More recently, Maas [4] considered Markov chains with finite state space and showed that, in this case, the Wasserstein distance does not allow this identification. Instead, assuming a time-reversible Markov chain, he was able to construct a different metric that allows this identification. Different constructions have been given in [1, 5] and the setting used is also described in [2].

The construction of the metric is involved and uses time-reversibility at several places. This further motivates our study of the converse of these statements. For this we will first introduce the setting used and define the gradient flow.

We consider continuous-time irreducible Markov chain with finite state space $\mathcal{X} = \{0, 1, \dots, d\}$. We denote its generator by $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ where for $i \neq j$ the entry Q_{ij} is the transition rate from state i to state j and $Q_{ii} = -\sum_{j\neq i} Q_{ij}$.

Given the initial probability distribution μ of the Markov chain, the probability distribution after time t will be given by μe^{tQ}. Note that the transition matrix e^{tQ} acts on the right on the row vector μ and the evolution of the law is captured by the Markov

*Support: UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/H023348/1 for the University of Cambridge Centre for Doctoral Training, the Cambridge Centre for Analysis.
†University of Cambridge, UK. E-mail: H.G.W.Dietert@maths.cam.ac.uk

semigroup e^{tQ}. Since the Markov chain is irreducible, there exists a unique stationary distribution π to which μe^{tQ} will converge as $t \to \infty$.

For the definition of a gradient flow, let $\mathcal{P}$ be the space of probability distributions on $\mathcal{X}$ with positive mass for any state. Then $\mathcal{P}$ can be naturally understood as the d-dimensional sub-manifold $\{v \in \mathbb{R}^{\mathcal{X}} : \sum_{i\in\mathcal{X}} v_i = 1 \text{ and } v_i > 0 \ \forall i \in \mathcal{X}\}$ of $\mathbb{R}^{\mathcal{X}}$. Under this identification, the tangent space at any point is

$$T = \{v \in \mathbb{R}^{\mathcal{X}} : \sum_{i\in\mathcal{X}} v_i = 0\}.$$

Given a functional $\mathfrak{F} : \mathcal{P} \mapsto \mathbb{R}$ and a Riemannian metric g on $\mathcal{P}$, the gradient flow $\rho : \mathbb{R}^+ \mapsto \mathcal{P}$ is determined by the differential equation

$$g|_{\rho(t)}(\dot{\rho}(t), v) = -\mathrm{d}\mathfrak{F}|_{\rho(t)}(v) \quad \forall v \in T, t \in \mathbb{R}^+.$$

That is $\dot{\rho}$ equals $-\mathrm{d}\mathfrak{F}$ under the identification of the tangent space and the cotangent space through the metric g.

We say that the gradient flow under g of $\mathfrak{F}$ equals the flow associated to the semigroup e^{tQ} if for all $\mu \in \mathcal{P}$ the trajectory $t \mapsto \mu e^{tQ}$ equals the gradient flow $\rho(t)$ with $\rho(0) = \mu$.

From the continuous setting a natural functional is the relative entropy $\mathfrak{H}$ with respect to the stationary state π, which is defined by

$$\mathfrak{H}(\rho) = -\sum_{i\in\mathcal{X}} \rho_i \log \frac{\pi_i}{\rho_i}.$$

Throughout this work the relative entropy is understood with respect to the stationary distribution of the considered Markov chain.

The result by Maas [4] now is: *There exists a metric g on $\mathcal{P}$ such that the gradient flow under g of the relative entropy $\mathfrak{H}$ equals the flow associated to the semigroup e^{tQ}.*

First we show that the metric is not unique:

Theorem 1.1. *For $d \geq 2$, consider a continuous-time irreducible Markov chain with finite state space $\mathcal{X} = \{0, 1, \ldots, d\}$, generator Q, and stationary distribution π. Let $\mathfrak{F} : \mathcal{P} \mapsto \mathbb{R}$ be a differentiable functional and g a Riemannian metric on $\mathcal{P}$. If the flow associated to the semigroup e^{tQ} equals the gradient flow of $\mathfrak{F}$ under the metric g, then for $\rho \in \mathcal{P}$ with $\rho \neq \pi$, there exists another metric $\tilde{g}$ on $\mathcal{P}$ such that $\tilde{g} \neq g$ at ρ and such that the gradient flow of $\mathfrak{F}$ under the metric $\tilde{g}$ still equals the flow associated to the semigroup e^{tQ}.*

As converse of the construction we show:

Theorem 1.2. *Consider a continuous-time irreducible Markov chain with finite state space $\mathcal{X} = \{0, 1, \ldots, d\}$, generator Q, and stationary distribution π. If the flow associated to the semigroup e^{tQ} equals the gradient flow of $\mathfrak{F} \in C^2$ under a Riemannian metric $g \in C^1$ on $\mathcal{P}$, then, with $g|_\pi$ as metric at π,*

(a) $g|_\pi$ is uniquely determined by Q and $\mathfrak{F}$,

(b) Q can be computed from $g|_\pi$ and $\mathfrak{F}$,

(c) Q is real diagonalisable,

(d) Q is time-reversible if $\mathfrak{F}$ is the relative entropy $\mathfrak{H}$.

Conversely, if Q is real diagonalisable, then there exists a Riemannian metric g and a smooth functional $\mathfrak{F}$ on $\mathcal{P}$ such that the gradient flow of $\mathfrak{F}$ under the metric g equals the flow associated to the semigroup e^{tQ}.

This result shows that the assumption of time-reversibility in the construction of the metric by Maas is necessary and cannot be relaxed. Moreover, the results for the generator Q come from the differentiability around the equilibrium distribution π, so that the theorem holds as long as there is a neighbourhood of π in which the Markov semigroup e^{tQ} equals the gradient flow.

Remark 1.3. *Theorems 1.1 and 1.2 are obtained by analysing the Riemannian structure of the gradient flow and thus can be formulated for general gradient flows on finite-dimensional manifolds. For this, part (d) of Theorem 1.2 can be formulated with a weighted ℓ^2-norm, i.e. Q is symmetric with respect to this norm, if $\mathfrak{F}$ is the squared distance to π under this ℓ^2-norm. In fact, relating the results to Q acting on the bigger space $\mathbb{R}^{\mathcal{X}}$ makes the proofs slightly longer.*

Combining this theorem with Maas' result gives our main theorem.

Theorem 1.4 (Characterisation of Markov chains)**.** *Consider a continuous-time irreducible Markov chain with finite state space and generator Q.*

- *The Markov chain is time-reversible if and only if there exists a metric g such that the flow associated to the semigroup e^{tQ} equals the gradient flow under $g \in C^1$ of the relative entropy with respect to the stationary distribution.*
- *The Markov chain has a real diagonalisable generator Q if and only if there exists a metric $g \in C^1$ and a functional $\mathfrak{F} \in C^2$ such that the flow associated to the semigroup e^{tQ} is the gradient flow of $\mathfrak{F}$ under g.*

The characterisation of real diagonalisable generators Q shows for example that Markov chains that also have an oscillatory behaviour cannot be described by gradient flows.

Finally, we remark that the relative entropy $\mathfrak{H}$ depends on the generator Q only through the equilibrium distribution π. Moreover, any functional $\mathfrak{F}$ that allows to construct a gradient flow for all time-reversible Markov chains must have a similar Taylor expansion around π. More precisely:

Theorem 1.5. *Fix a finite state space $\mathcal{X}$, a distribution $\pi \in \mathcal{P}$ and a functional $\mathfrak{F} : \mathcal{P} \mapsto \mathbb{R}$ in C^2. Suppose that, for every generator Q defining an irreducible time-reversible Markov chain with state space $\mathcal{X}$ and stationary distribution π, there exists a Riemannian metric $g \in C^1$ on $\mathcal{P}$ such that the gradient flow of $\mathfrak{F}$ under the metric g equals the flow associated to the semigroup e^{tQ}. Then there exists a positive constant α such that*

$$\mathrm{d}\mathfrak{F}|_\pi = \mathrm{d}\mathfrak{H}|_\pi = 0 \text{ and } \mathrm{d}^2\mathfrak{F}|_\pi = \alpha\, \mathrm{d}^2\mathfrak{H}|_\pi,$$

where $\mathfrak{H}$ is the relative entropy with respect to π.

Here we use the notation $\mathrm{d}\mathfrak{F}|_\pi$ to denote the first derivative of $\mathfrak{F}$ at the point π, which we understand as linear map from T to $\mathbb{R}$. With $\mathrm{d}^2\mathfrak{F}|_\pi$ we denote the second derivative at the point π, which is a linear map from $T \times T$ to $\mathbb{R}$.

The assumption on the functional $\mathfrak{F}$ is not empty, because Maas' result states that the relative entropy $\mathfrak{H}$ with respect to π is a functional satisfying the assumption. Moreover, it cannot be strengthened to uniqueness. For this, another functional is the quadratic form $\mathfrak{F}$ defined by $\mathfrak{F}|_\pi = \mathrm{d}\mathfrak{F}|_\pi = 0$ and $\mathrm{d}^2\mathfrak{F}|_\pi = \alpha \mathrm{d}^2\mathfrak{H}|_\pi$ for some $\alpha > 0$. This satisfies the assumption, because for any generator Q the constant metric g given by the value $g|_\pi$ in part (a) of Theorem 1.2 indeed defines a Riemannian metric with the required identification.

Remark 1.6. *In [4], Maas considered continuous-time Markov chains obtained from an irreducible discrete-time Markov chain with transition matrix K by choosing the jump times according to a Poisson process. The resulting generator is $Q = K - I$ and, by analogy with continuous-time diffusion processes, the semigroup e^{tQ} is also called a heat flow.*

By time-rescaling, this is no restriction to the class of Markov chains for the study of gradient flows because, for any generator Q, we can find some $\alpha > 0$ such that $K = I + \alpha Q$ is non-negative along the diagonal and K defines a transition matrix. Now given a functional $\mathfrak{F}$, if we can find a metric g such that the flow associated to the semigroup $e^{t(\alpha Q)}$ equals the gradient flow of $\mathfrak{F}$ under g, then the flow associated to the semigroup e^{tQ} equals the gradient flow of $\mathfrak{F}$ under the rescaled metric g/α.

2 Characterisation of Markov chains

Recall that for an irreducible Markov chain the evolution μe^{tQ} converges to the unique stationary distribution π for any $\mu \in \mathcal{P}$. Hence μQ vanishes if and only if $\mu = \pi$. With this observation, we can construct a perturbation of the metric in order to prove Theorem 1.1.

Proof of Theorem 1.1. Let e_1 be the vector field given by μQ at $\mu \in \mathcal{P}$. For a small enough neighbourhood $V \subset \mathcal{P}$ of ρ, we can find smooth vector fields $e_2, \dots, e_d$ such that $e_1, \dots, e_d$ is a basis of T at every $\mu \in V$. Let η be a smooth cutoff functional with compact support, vanishing outside V, and satisfying $\eta(\rho) \neq 0$. Then define another metric $\tilde{g}$ by

$$\tilde{g}|_\mu(e_i, e_j) = \begin{cases} g(e_i, e_j)|_\mu + \eta(\mu)a & \text{if } i = j = 2, \\ g(e_i, e_j)|_\mu & \text{otherwise,} \end{cases}$$

for $\mu \in V$ and a constant $a \in \mathbb{R}$. Outside of V, define $\tilde{g} = g$.

If a is small enough, $\tilde{g}$ is still positive definite and is therefore a metric. Moreover, $\tilde{g}$ creates the same gradient flow, because for every $\mu \in \mathcal{P}$ and any $v \in T$

$$\tilde{g}|_\mu(\mu Q, v) = \tilde{g}|_\mu(e_1, v) = g|_\mu(e_1, v) = g|_\mu(\mu Q, v). \qquad \square$$

For the characterisation at the equilibrium distribution, we use the assumed differentiability of g and $\mathfrak{F}$.

Proof of Theorem 1.2. The equality of the flow associated to the semigroup e^{tQ} and the gradient flow implies that the time derivatives of both evolutions agree at every state $\pi + h$ with $h \in T$. This means that, for all $v \in T$,

$$g|_{\pi+h}((\pi + h)Q, v) = -\mathrm{d}\mathfrak{F}|_{\pi+h}(v). \tag{2.1}$$

Since $\pi Q = 0$, this implies that $\mathrm{d}\mathfrak{F}$ must vanish at π. Moreover, it simplifies Equation (2.1) to

$$g|_{\pi+h}(hQ, v) = -\mathrm{d}\mathfrak{F}|_{\pi+h}(v).$$

Let $M = \mathrm{d}^2\mathfrak{F}|_\pi$, then the Taylor expansion around $h = 0$ shows by the assumed regularity of g and $\mathfrak{F}$ that

$$g|_\pi(hQ, v) + O(\|h\|^2) = -M(h, v) + O(\|h\|^2).$$

As this holds for arbitrary $h \in T$, the linear terms must agree. Hence,

$$g|_\pi(wQ, v) = -M(w, v) \quad \forall v, w \in T. \tag{2.2}$$

Furthermore, we claim that the restriction of Q to T defines an automorphism on T. Since Q preserves the probability mass (i.e. $\sum_{i \in \mathcal{X}} Q_{ji} = 0$ for $j \in \mathcal{X}$), its range is

inside T. If Q was not an automorphism, a $v \in T$ satisfying $vQ = 0$ would exist by the Rank-Nullity Theorem. But then, for small enough α, also $\pi + \alpha v$ would be a stationary state, contradicting the irreducibility of the Markov chain.

Hence Equation (2.2) determines the value of $g|_\pi$ for all arguments, which proves part (a) of the theorem.

Given $g|_\pi$ and M, Equation (2.2) determines $vQ \cdot w$ for all $v, w \in T$, because $g|_\pi$ is a positive form. By the mass conservation $vQ \in T$, so that this determines vQ for all $v \in T$. Since $\pi Q = 0$, this determines Q and shows part (b).

In order to prove part (c) let $f_1, \dots, f_d$ be a basis of T which is orthonormal under $g|_\pi$, i.e. $g|_\pi(f_i, f_j) = \delta_{ij}$. Let $\bar{Q}$ be the matrix corresponding to the generator Q in this basis, i.e. $f_i Q = \sum_{j=1}^d \bar{Q}_{ij} f_j$ for $i = 1, \dots, d$. Also let $\bar{M}$ be the matrix corresponding to M in this basis, i.e. $\bar{M}_{ij} = M(f_i, f_j)$. Then Equation (2.2) becomes

$$x\bar{Q} \cdot y = -x\bar{M} \cdot y \quad \forall x, y \in \mathbb{R}^d.$$

Hence $\bar{Q} = -\bar{M}$. Since the partial derivatives of $\mathfrak{F}$ commute, $\bar{M}$ is symmetric. Therefore, $\bar{Q}$ is real diagonalisable and $\bar{Q}$ has d real eigenvectors in T. Since π is another eigenvector of Q not in T, this implies that Q is real diagonalisable, which is the statement of part (c).

For the converse of the theorem, we assume that Q is diagonalisable and we need to construct a suitable functional and metric on $\mathcal{P}$. For this, fix eigenvectors $\pi, f_1, \dots, f_d$ of Q with eigenvalues $0, \lambda_1, \dots, \lambda_d$. Since π is the only stationary distribution, $\lambda_i \neq 0$ for $i = 1, \dots, d$. As Q maps into T, this shows that $f_1, \dots, f_d$ is a basis of T. Define the constant Riemannian metric g on $\mathcal{P}$ by

$$g(f_i, f_j) = \delta_{ij},$$

and the functional $\mathfrak{F} : \mathcal{P} \mapsto \mathbb{R}$ by

$$\mathfrak{F}(\pi + \sum_{i=1}^d a_i f_i) = \frac{1}{2} \sum_{i=1}^d (-\lambda_i) a_i^2.$$

Then at any state $\mu = \pi + \sum_{i=1}^d a_i f_i \in \mathcal{P}$ we have, for $j = 1, \dots, d$,

$$g(\mu Q, f_j) = \lambda_j a_j \text{ and } -\mathrm{d}\mathfrak{F}|_\mu(f_j) = \lambda_j a_j.$$

Hence the flow associated to the semigroup e^{tQ} and the gradient flow agree, because their time derivatives agree for every probability distribution $\mu \in \mathcal{P}$.

For the remaining part (d), the functional $\mathfrak{F}$ is the relative entropy $\mathfrak{H}$. The second derivative $\mathrm{d}^2\mathfrak{H}|_\pi$ of $\mathfrak{H}$ at π is given by

$$M(w, v) = \sum_{\alpha \in \mathcal{X}} \frac{w_\alpha v_\alpha}{\pi_\alpha}$$

for $v, w \in T$. Let $\Pi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be the diagonal matrix with diagonal entries $(\pi)_{i \in \mathcal{X}}$. Then, by the calculated form of M, we have $v\Pi^{-1} \cdot w = M(v, w)$ for all $v, w \in T$.

Over T, the metric g has an inverse b at π which is defined by $g|_\pi(v, wb) = v \cdot w$ for $v, w \in T$ and is a positive definite symmetric automorphism on T. Define the symmetric matrix $a \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ by $va = vb$ for $v \in T$ and $\mathbf{1}a = 0$, where $\mathbf{1}$ is the vector with all entries 1.

Since b is an automorphism on T, Equation (2.2) implies $g(vQ, ua) = -v\Pi^{-1} \cdot (ua)$ for all $u, v \in T$. By the construction of a, this shows that for all $u, v \in T$

$$vQ \cdot u = -v\Pi^{-1} a \cdot u. \tag{2.3}$$

Since πQ and $\pi\Pi^{-1}a$ both vanish, Equation (2.3) also holds for $v = \pi$. Hence it holds for all $v \in \mathbb{R}^{\mathcal{X}}$ which shows $Q \cdot u = -\Pi^{-1}a \cdot u$ for all $u \in T$.

By the conservation of probability $Q \cdot \mathbf{1} = 0$ and by the symmetry of a also $a \cdot \mathbf{1} = 0$, so that $Q \cdot u = -\Pi^{-1}a \cdot u$ holds for all $u \in \mathbb{R}^{\mathcal{X}}$. This shows

$$\Pi Q = -a.$$

Since a is symmetric, this shows that Q satisfies the detailed balance equation, i.e. that the Markov chain is time-reversible. □

The remaining Theorem 1.5 is reduced to the following lemma.

Lemma 2.1. *Assume the hypothesis of Theorem 1.5. Then $\mathfrak{F}$ and $\mathfrak{H}$ have a minimum at π, and the derivatives $M = \mathrm{d}^2\mathfrak{F}|_\pi$ and $N = \mathrm{d}^2\mathfrak{H}|_\pi$ are positive non-degenerate forms on T. For every $v \in T$ vanishing in exactly one state, there exists $\alpha_v \in \mathbb{R}^+$ such that $M(v,\cdot) = \alpha_v N(v,\cdot)$.*

With this lemma the theorem can be proved.

Proof of Theorem 1.5. By the previous lemma $\mathfrak{F}$ and $\mathfrak{H}$ have a minimum at π so that $\mathrm{d}\mathfrak{F}_\pi = \mathrm{d}\mathfrak{H}_\pi = 0$.

If $d = 1$, then M and N from the lemma correspond to positive scalars so that there exists a positive scalar α satisfying the required relation $M = \alpha N$.

Hence, assume $d \geq 2$. For two states $v, \tilde{v} \in T$ vanishing only in one common state, the constants α_v and $\alpha_{\tilde{v}}$ of the lemma must agree. If v and $\tilde{v}$ are linearly dependent, then this follows directly from the bilinearity of M and N. Otherwise, for small enough $\lambda \in \mathbb{R}$, also $v + \lambda\tilde{v}$ is in T and vanishing in exactly one state. Hence the lemma applies to this state and by linearity follows

$$\alpha_{v+\lambda\tilde{v}} N(v + \lambda\tilde{v}, \cdot) = \alpha_v N(v,\cdot) + \lambda\alpha_{\tilde{v}} N(\tilde{v},\cdot).$$

Since N is non-degenerate, the functionals $N(v,\cdot)$ and $N(\tilde{v},\cdot)$ are linearly independent, so that the equation implies $\alpha_{v+\lambda\tilde{v}} = \alpha_v$ and $\alpha_{v+\lambda\tilde{v}} = \alpha_{\tilde{v}}$ and thus, as claimed, $\alpha_v = \alpha_{\tilde{v}}$.

Hence, for $i \in \mathcal{X}$, there exists $\bar{\alpha}_i$ such that $M(v,\cdot) = \bar{\alpha}_i N(v,\cdot)$ holds for $v \in T$ with $v_i = 0$ and $v_j \neq 0$ for $j \neq i$. By continuity of $M(v,\cdot)$ and $\bar{\alpha}_i N(v,\cdot)$ with respect to v, the result also holds for all $v \in T$ with $v_i = 0$.

If $d \geq 3$ then for $i, j \in \mathcal{X}$ there exists $v \in T$ with $v_i = v_j = 0$ which implies that all $\bar{\alpha}_i$ for $i \in \mathcal{X}$ agree. If $d = 2$, we can consider $v = (1\ {-1}\ 0)$ and $\tilde{v} = (0\ 1\ {-1})$ with $\alpha_v = \bar{\alpha}_2$ and $\alpha_{\tilde{v}} = \bar{\alpha}_0$. Then $v + \tilde{v}$ satisfies the lemma as well so that as before $\bar{\alpha}_0 = \bar{\alpha}_2$. Likewise, $\bar{\alpha}_0 = \bar{\alpha}_1$ and all $\bar{\alpha}_i$ agree.

The common value α of the $\bar{\alpha}_i$ for $i \in \mathcal{X}$ is the claimed constant satisfying $\mathrm{d}^2\mathfrak{F}|_\pi = \alpha\mathrm{d}^2\mathfrak{H}|_\pi$. Since the derivatives $\mathrm{d}^2\mathfrak{F}|_\pi$ and $\mathrm{d}^2\mathfrak{H}|_\pi$ are positive forms, the constant α must be positive. □

The remaining lemma is proved by considering suitable Markov chains.

Proof of Theorem 2.1. Since π has positive mass for every state $i \in \mathcal{X}$, there exists a time-reversible irreducible Markov chain with stationary state π. Consider such a Markov chain with generator Q and let g be a metric such that the flow associated to the Markov semigroup e^{tQ} is the gradient flow of $\mathfrak{F}$ under g. Then for any initial probability state $\mu \in \mathcal{P}$ the value $\mathfrak{F}(\mu e^{tQ})$ is decaying as t increases. Hence $\mathfrak{F}$ must have a minimum at π. As in the proof of Theorem 1.4, Equation (2.2) must hold and $g|_\pi$ and Q are non-degenerated over T. Hence $M = \mathrm{d}^2\mathfrak{F}|_\pi$ must be non-degenerated and positive, because $\mathfrak{F}$ has a minimum at π. Since $\mathfrak{H}$ satisfies the assumptions imposed on $\mathfrak{F}$, this also holds for $\mathfrak{H}$.

If $d = 1$, there exists no such state $v \in T$. Thus assume $d \geq 2$ henceforth. Moreover, label the states such that $v_d = 0$ and identify M, N and $g|_\pi$ with the corresponding matrix, i.e. $M(v, w) = vM \cdot w$, and $N(v, w) = vN \cdot w$, and $g|_\pi(v, w) = vg|_\pi \cdot w$ for all $v, w \in T$.

Then, Equation (2.2) implies $M^{-1}Q = -g|_\pi^{-1}$. Hence $M^{-1}Q$ is symmetric and thus $M^{-1}QM = Q^T$. Likewise $N^{-1}QN = Q^T$ and thus $MN^{-1}QNM^{-1} = Q$. Therefore, if v is an eigenvector of Q, then vMN^{-1} is again an eigenvector of Q with the same eigenvalue.

We finish the proof of the lemma by showing that v and $w := vMN^{-1}$ are proportional. For this, we show $w_d = 0$ and $w_i v_j = v_i w_j$ for all $i, j = 0, 1, \ldots, d-1$ by constructing suitable Markov chains whose generator Q has the eigenvector v.

For the first part, let $\lambda > 0$ and consider the time-reversible irreducible Markov chain with stationary state π and generator matrix

$$Q = \Pi^{-1} \begin{pmatrix} -\pi_0\lambda & 0 & 0 & \ldots & 0 & \pi_0\lambda \\ 0 & -\pi_1\lambda & 0 & \ldots & 0 & \pi_1\lambda \\ 0 & 0 & -\pi_2\lambda & \ldots & 0 & \pi_2\lambda \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & -\pi_{d-1}\lambda & \pi_{d-1}\lambda \\ \pi_0\lambda & \pi_1\lambda & \pi_3\lambda & \ldots & \pi_{d-1}\lambda & -(1-\pi_d)\lambda \end{pmatrix}.$$

The eigenspace for $-\lambda$ is $\{u \in T : u_d = 0\}$ and contains v. Hence, $w_d = 0$.

By further relabelling the states, it suffices to show for the second case $w_0 v_1 = v_0 w_1$. For this, let $\lambda > 0$ and consider the Markov chain with generator

$$Q = \Pi^{-1} \begin{pmatrix} -\pi_0\lambda - \beta_0 & \mu & 0 & \ldots & 0 & \pi_0\lambda + \beta_0 - \mu \\ \mu & -\pi_1\lambda - \beta_1 & 0 & \ldots & 0 & \pi_1\lambda + \beta_1 - \mu \\ 0 & 0 & -\pi_2\lambda & \ldots & 0 & \pi_2\lambda \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & -\pi_{d-1}\lambda & \pi_{d-1}\lambda \\ \pi_0\lambda + \beta_0 - \mu & \pi_1\lambda + \beta_1 - \mu & \pi_2\lambda & \ldots & \pi_{d-1}\lambda & \gamma \end{pmatrix}$$

where

$$\beta_0 = \mu\frac{\pi_0 v_1}{v_0 \pi_1}, \qquad \beta_1 = \mu\frac{\pi_1 v_0}{v_1 \pi_0}, \qquad \gamma = -(1-\pi_d)\lambda - \mu\frac{\pi_0 v_1}{v_0 \pi_1} - \mu\frac{\pi_1 v_0}{v_1 \pi_0} + 2\mu.$$

By choosing $\mu > 0$ small enough, this defines an irreducible time-reversible Markov chain with stationary state π and $vQ = -\lambda v$. Hence, w is again an eigenvector with eigenvalue $-\lambda$. From the first component, we find the required ratio

$$-\lambda w_0 = -\lambda w_0 - \mu\frac{v_1 w_0}{v_0 \pi_1} + \mu\frac{w_1}{\pi_1} \Rightarrow v_1 w_0 = v_0 w_1. \qquad \square$$

References

[1] Shui-Nee Chow, Wen Huang, Yao Li, and Haomin Zhou, *Fokker-Planck equations for a free energy functional or Markov process on a graph*, Arch. Ration. Mech. Anal. **203** (2012), no. 3, 969–1008. MR-2928139

[2] Matthias Erbar and Jan Maas, *Gradient flow structures for discrete porous medium equations*, Discrete Contin. Dyn. Syst. **34** (2014), no. 4, 1355–1374. MR-3117845

[3] Richard Jordan, David Kinderlehrer, and Felix Otto, *The variational formulation of the Fokker-Planck equation*, SIAM J. Math. Anal. **29** (1998), no. 1, 1–17. MR-1617171

[4] Jan Maas, *Gradient flows of the entropy for finite Markov chains*, J. Funct. Anal. **261** (2011), no. 8, 2250–2292. MR-2824578

[5] Alexander Mielke, *A gradient structure for reaction-diffusion systems and for energy-drift-diffusion systems*, Nonlinearity **24** (2011), no. 4, 1329–1346. MR-2776123

[6] Cédric Villani, *Optimal transport*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 338, Springer-Verlag, Berlin, 2009, Old and new. MR-2459454

Acknowledgments. The author would like to thank James Norris for many helpful discussions and suggestions.